Uber: Utilizing Buffers to Simplify NoCs for Hundreds-Cores

Giorgos Passas
gpassas81@gmail.com

Abstract—Approaching ideal wire latency using a network-on-chip (NoC) is an important practical problem for many-core systems, particularly hundreds-cores. Although other researchers have focused on optimizing large meshes, bypassing or speculating router pipelines, or creating more intricate logarithmic topologies, this paper proposes a balanced combination that trades queue buffers for simplicity. Preliminary analysis of nine benchmarks from PARSEC and SPLASH using execution-driven simulation shows that utilization rises from 2% when connecting a single core per mesh port to at least 50%, as slack for delay in concentrator and router queues is around 6× higher compared to the ideal latency of just 20 cycles. That is, a 16-port mesh suffices because queueing is the uncommon case for system performance. In this way, the mesh hop count is bounded to three, as load becomes uniform via extended concentration, and ideal latency is approached using conventional four-stage pipelines for the mesh routers together with minor logarithmic edges. A realistic *Uber* is also detailed, featuring the same performance as a 64-port mesh that employs optimized router pipelines, improving the baseline by 12%. Ongoing work develops techniques to better balance load by tuning the placement of cache blocks, and compares *Uber* with bufferless routing.

arXiv:1607.07766v2 [cs.AR] 27 Jul 2016

1 Introduction

To efficiently manage the high wire-to-gate latency ratios in modern VLSI, one popular processor architecture partitions the chip into many identical cores processing in parallel and communicating implicitly via loads and stores to a shared memory that is distributed with the cores [14], [18]. In such systems, a communication medium that uses core-to-core links would imply unmanageable, spaghetti wiring. Thus, together with the processor and memory slice each core also contains a router, and routers are connected in a regular topology, or network-on-chip (NoC). Moreover, NoCs are usually meshes for simplicity. Still, mesh routers are non-negligible overheads, particularly in latency. That being the case, and given the dependence of system performance [16], approaching ideal wire latency using a mesh NoC is an important practical problem [16], [17], [20], [21], [23]. Moreover, this problem becomes compounded as systems scale [1], [16]. Solutions include techniques to bypass routers [16], [17], [23] and low-latency router design [20], [21]. However, such techniques increase design complexity, while router overheads might remain high. What is more, wide links and high frequencies for fast serialization of core messages have led to such low utilizations [4], [10], [11], [13] that bufferless routing has been defended [19]. Taking into consideration all the above data, this paper proposes to compensate for system scaling by extending core concentration. Although delays in concentrator and router queues increase as mesh port load rises, queueing is the uncommon case for system performance. Besides, concentrators are local structures, and more intricate, globally logarithmic topologies [3] are obviated. In this way, this paper makes a clear case for a buffered NoC.

In particular, from present small systems to hundreds-cores of the near research future, mesh hop count makes a critical step, while port utilization does not exceed a typical low of 2%. Concentrating few cores is usual [3], but non scalable. This paper extends concentration to 16 and beyond, thus increasing utilization to 25% or higher. The utilization boost is enabled by state-of-the-art benchmarks that tolerate queue delays at least up to 3× higher compared to an ideal latency of just 20 cycles. Owing to this slack, a 16-port mesh scales to hundreds-cores, and hop count is easily bounded to three. Overall, the traditional 4-stage pipeline [8] suffices without any bypassing or speculation optimizations, and buffers implementing queues play a key role in the simplification. Although some researchers have already studied evenly-utilized configurations [7], [13], or even similar topologies [9], their analysis focuses on small systems or custom interconnects, hence missing the role of buffers.

Compatibly with previous studies (e.g. [13]), numbers correspond to analysis of seven benchmarks from PARSEC [5] and two benchmarks from SPLASH [28] on 64-cores in the gem5 full-system simulator [6]. In this context, this paper measures a cumulative load of 96% on average, and spreads this load to four ports. Assuming that (i) the need for bandwidth grows linearly with system size [1], and (ii) benchmarks for hundreds-core should demonstrate similar communication patterns like PARSEC, such a 4-port mesh is miniature of a 16-port counterpart in 256-cores. To better comprehend stress, the analysis abstracts the mesh using a single-stage switch that implements ideal output queueing [15]. As a transient step to reality, output queueing is replaced first by models of idealistic single-stage crossbars [15], [24]. In comparison, using a single FIFO per input, queue delay slips by head-of-line (HoL) blocking [15] slightly beyond the nominal slack, resulting to tangible system slowdown that measures to 1.08. Using more advanced organizations [22], queueing is sufficiently bounded, while scheduling efficiency plays a marginal role. Furthermore, because extended concentration makes the load uniform and bursty [25] crossbars are redundant, and are replaced by a mesh. Performance remains excellent, owing to router buffers that provide a kind of speedup. Comparing a 4-port and a 64-port mesh, both using 4-stage routers, the small instance reduces performance loss from 12% to 8%, although it presents higher end delay by queueing, and falls short of a large instance that offers ideal performance using single-cycle routers. This handicap is attributed to uncontrolled interleaving of control and data cells. Indeed, a more realistic organization that separates messages in virtual networks removes the above handicap. What is more, virtualization roughly doubles the slack for queueing.

TABLE 1
Additional System Parameters

PROCESSOR	2GHz, in order
L1 CACHE	(64+32)KB, private, 2 ways, 3 cycles
L2 CACHE	2MB/core [19], shared, 8 ways, 12 cycles, MESI full-map dir
MEMORY	100 cycles, 72B blocks
SPLASH	258×258 matrix, 1M integers
OS	Linux 2.6.27, 1 thread/core

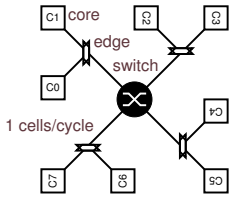

Fig. 1. NoC model diagram

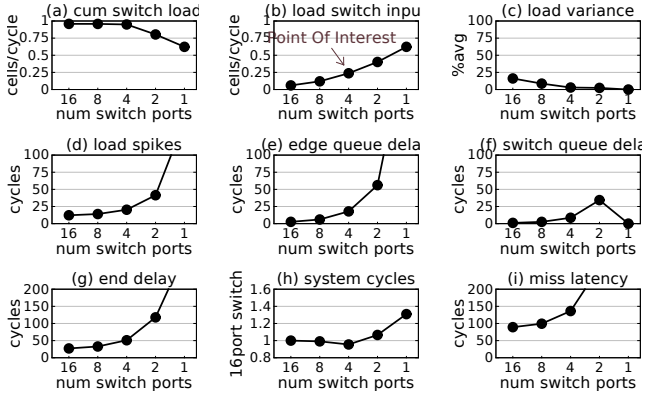

Fig. 2. (a-g) NoC stress and (h-i) system performance; each point is the average of nine benchmarks

My Contributions:

- Analysis to compare the latency and queue delay impact of a NoC on system performance (Sec. 3, Sec. 4)
- Analysis to plot HoL blocking as a side-effect of utilization boost (Sec. 3, Sec. 5)
- A highly utilized 16-port mesh for hundreds-cores featuring simpler pipelines by better utilized buffers (Sec. 6, Sec. 7)

2 System and NoC Model

The focus is on a concrete memory organization, similar as the baseline in [18]. Memory blocks are 72 B and coherence control messages are 8 B. More details are given in Table 1. With respect to benchmarks, this paper considers PARSEC [5] and SPLASH [28] for comparison with PARSEC. CPI is on average seven, when measured on 4-cores. For experiments using unloaded NoCs, I was always able to provide five runs for all benchmarks except for streamcluster (one run) and facesim (two runs). With very stressed NoCs, runs are in general fewer. The critical path is ≈10 days and the typical case is ≈4 days. Finally, for PARSEC, input sets are simlarge, whereas for SPLASH inputs correspond to simsmall sets.

Fig. 1 gives a diagram of the NoC. Core messages are fragmented into fixed-size cells, which cells are injected through edge multiplexors, switched, and ejected through the demultiplexors. Once messages are reassembled from cells at their destination core, they are delivered to the correct controller. The term *cell* emphasizes that queues are primarily infinite. In particular, the switch employs one queue per output so that cells that arrive concurrently at the switch inputs are written in parallel and indiscriminately. Such output queueing [15] is used widely for ideal-performance switching. Output queueing is also useful to model random multiplexing at the edges. Note that contending cells are always downstream, and there is no contention from edges to cores. Following the above discussion, although the switch is ideal, the whole NoC is not. Referring to Fig. 1, consider three cells at cores C0, C1, and C2 destined to C4, C6, and C6, respectively (C6 is double to denote contention). When C1 is "accidentally" prioritized over C0 at the edge, the NoC suffers one extra delay cycle. Nonetheless, I was unable to plot any performance loss by such edge blocking. The pipeline comprises 3 cycles at the links, plus 1 cycle at every other component (including the cores). By default links are 4 B instead of a typical 16 B [13] to reserve about 4× bandwidth for faster systems and more demanding benchmarks. Thus, adding 10 cycles for message serialization, the total latency is 26 cycles. End delay is offset by queueing and message reassembly.

3 Evaluations Using Idealistic NoCs

This section evaluates NoC stress and the impact thereof on system performance, focusing on 4-port switches in 64-cores as miniatures of 16-port switches for 256-cores. The main variable is the size of the switch, which ranges from 16 ports down to one port. A NoC using a 1-port switch is actually a bus [7], [27].

Thus, Fig. 2 plots nine metrics averaged over the nine benchmarks of PARSEC and SPLASH. In (a), only a single link is fully utilized on average. Such a low utilization is well known in the literature [4], [10], [11], [13]. Moreover, load drops using small switches. System slowdown largely explains this throttling. In (b), cumulative load is averaged over the switch inputs. We observe that load remains below 1 cells/cycle. The number of interest is 0.24 cells/cycle for 4-port switches. In (c), switch inputs are more uniformly loaded for smaller switches. In such a way, edges never saturate. Indeed, edge queueing is plotted comparatively short below. A similar discussion applies also for the outputs of the switch, hence neither does the switch saturate. In (d), we observe that switch load is spiked. Spikes start from 10 cycles —corresponding to roughly one control message every other cache block— growing up to several tens of cells. For 4-port switches, spikes are 20 cycles. Although NoC load patterns are known to contain spikes [4], [11], this paper is from the first to clearly plot their impact on system performance, in a way merging spikes in time and/or space. Overall, switch loading resembles the Internet core [25]. Load metrics are helpful for understanding NoC behavior, but what matters for the system is end delay. Fig. 2(e) and (f) contribute queue delays at the edges and the switch, respectively. We observe that queue delays increase as the switch shrinks. Using 4-port switches, 18 cycles are spent at the edges and 9 cycles at the switch. This 2× ratio is consistent also for smaller switches. Further than queue delays, the time it takes to reassemble messages from cells at the cores increases. In any case reassembly delay is almost negligible (not shown). Next, (g) plots end message delay. Using 4-port switches, end delay is on average 53 cycles, of which 26 cycles is the latency term and the remainder is the sum of queue delays at the edges and the switch. Note, we assumed that data and control messages suffer equal delays in the NoC. This paper found the above assumption true, owing to NoC multiplexing being random. Finally, the impact of queue delays on system performance is plotted in Fig. 2(h). Shrinking the switch down to 4-port is no trouble. Slowdown starts using 2-port switches; and measures to 1.07 in this case. Next, Fig. 2(i) plots the L1 miss latency, as reported in the system statistics [6]. Using this first-order approximation,

$$\text{Miss Latency} = L_{L1} + D_{ee}^{ctrl} + L_{L2} + D_{ee}^{data} + \text{L2 Miss Rate} \times D_x,$$

hits in L2 explain 75% of the miss latency.

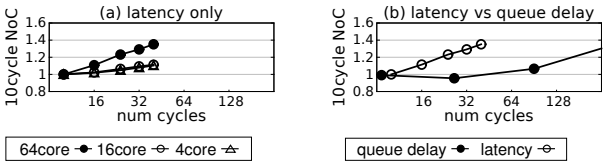

Fig. 3. Effect of latency and queue delay on system performance; latency is total, end-to-end; queue delay is total, at the edges and the switch

4 Queue Delay Compared to Latency

This section adds results from a complementary experiment that varies NoC latency by explicitly offsetting the links. Fig. 3(a) plots system performance under ranging NoC latency for three distinct system sizes. We observe that 64-cores are more sensitive than smaller systems. This differentiation is due to a single benchmark, namely streamcluster. Either way, 40 cycles imply system slowdown of at least 1.09. Though marginal, slowdown is clear even at 20 cycles. Although Krishna et al. [16] measure significant losses at even lower latencies, Enright-Jerger et al. [10] are more conservative, as is this paper. Fig. 3(b) compares latency and queue delay. The queue delay curve merges the sum of edge and switch delays in Fig. 2(e) and (f) with the corresponding system performance in Fig. 2(h). Thus, each data point refers to NoC using a distinct switch instance from 8-port down to 1-port. In comparison, queue delay can grow 3× higher. This experimental result is intuitive: both inside each individual benchmark and across the whole set, communication can be seen in phases where load is typically low [4].

5 Replacing Ideal Switches with Crossbars

This section evaluates system performance degradation when ideal output queueing is replaced with more practical crossbars. This transition is particularly interesting given recent studies that have demonstrated the feasibility of large crossbars on chip [24].

A first model of input queueing (IQ) employs a single FIFO queue per input together with one round-robin arbiter per crossbar output [15]. Fig. 4 compares IQ with OQ. We observe in (a) that switch load is slightly lower in IQ. Edges reach equilibrium with the switch as benchmarks slow down by HoL blocking [15] at the switch, despite the fact that there is no backpressure (unlimited queues). In turn, edge queueing is also shorter as in Fig. 4(b). However, switch queueing is longer by HoL blocking as in (c). For 4-port switches, delay grows from 9 cycles in OQ to 51 cycles in IQ. Thus, delay is longer overall, and benchmarks slow down as in Fig. 4(d). For 4-port switches, slowdown measures to 1.08. Note that the above delay-performance relation is a good match with the analysis in Sec. 4.

There is a large body of research on improving IQ. A popular solution is Virtual Output Queueing together with a crossbar scheduler like *i*SLIP [22]. Fig. 5 plots system performance separately for each benchmark for 4-port and 2-port switches. We observe that bodytrack, canneal, and facesim are critical benchmarks that differentiate scheduling performance. On average, however, differences are impractical. Moreover, observe in Fig. 5 that there is an anomaly in the performance of canneal with 4SLIP, despite the fact that 4SLIP queueing is actually as good on average. This anomaly returns in Sec. 6 for the mesh. The corner behavior of RRM [22], on the other hand, is well explained by

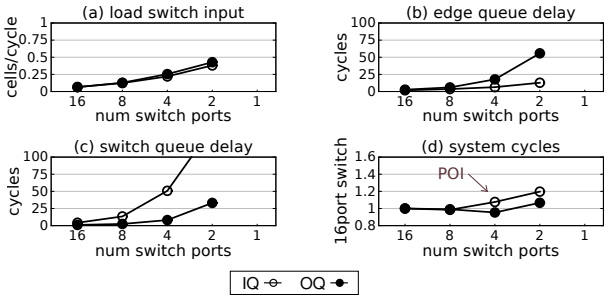

Fig. 4. Effect of HoL blocking on NoC and system performance

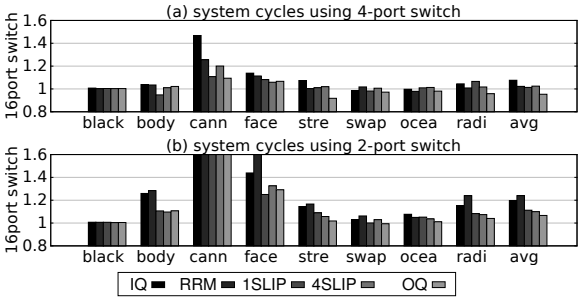
(a) system cycles using 4-port switch
(b) system cycles using 2-port switch
16port switch
black body cann face stre swap ocea radi avg
IQ RRM 1SLIP 4SLIP OQ

Fig. 5. Crossbar performance per benchmark; fluidanimate failed for some configurations

synthetic-load experiments. In conclusion, HoL blocking and its efficient reduction is a basic problem associated with utilization boost. Either way, the role of techniques to resolve HOL blocking is corrective, to prevent queueing from slipping beyond a nominal slack.

6 Replacing Crossbars with a Mesh

This section details the transition from crossbar to mesh. A mesh provides a finer physical hierarchy and also simplifies queueing. Parameters of interest are as follows: Routers are always dimension-ordered, router queues infinite FIFOs at inputs, router latency 4-stage, and link latency one cycle. For small instances, mesh hop count is bounded to three, hence latency is as for the crossbar. For 64-port instances, however, hop count jumps to seven, and latency makes a critical step.

Fig. 6 compares mesh and crossbar queue delay under a synthetic load. The mesh is as good as 1SLIP, shifting toward IQ for larger instances (not shown). Such a nice behavior for small instances should be owing to router buffers that provide a kind of speedup. The performance under benchmark load is plotted in Fig. 7. The plot intentionally drops ocean and streamcluster to reduce both noise and bias in favor of latency; this work-around could be tolerated considering ocean as a good approximation in place of streamcluster. The mesh is almost as good as expected. The handicap that appears for canneal and ocean is attributed to cells of control messages suffering exceptionally higher delays. Furthermore, for both the crossbar and the mesh, the benchmark load gives the same delays as double the synthetic load. This discrepancy should be due to spikes being roughly double than bursts by serialization. Unfortunately, this paper had no capacity to study correlation further than spike duration. A complementary study is [4].

Fig. 8 compares a 4-port mesh and a 64-port mesh. The 4-port mesh is better, although the end delay is higher by queueing. In

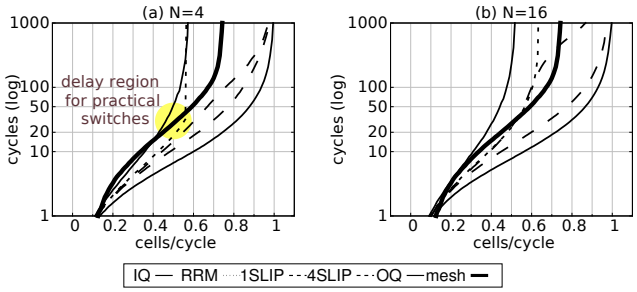

Fig. 6. Mesh compared to crossbar under synthetic load; links are 4B

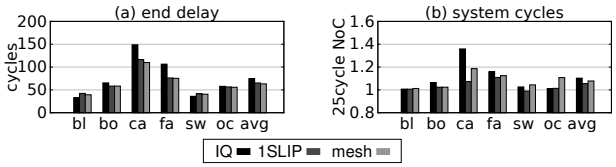

Fig. 7. Mesh compared to crossbar under benchmark load (execution driven); all instances are 4-port; streamcluster and radix are dropped

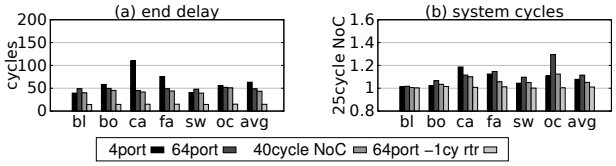

Fig. 8. Comparing 4-port and 64-port meshes; large meshes use either 4-cycle or 1-cycle routers; uniform latency of 40 cycles is also included

particular, the 64-port mesh degrades performance by a clear 12%. Moreover, performance is worse when compared to a uniform latency of 40 cycles. This is why researchers are proposing techniques to improve locality by optimizing placement of cache blocks [9], [27]. A third bar in Fig. 8, corresponding to a 64-port mesh using one-cycle routers, represents router optimization techniques [17], [20], [21], [23]. Clearly, a large mesh outperforms a small mesh when such techniques are employed. However, the small mesh can be improved by fixing the interleaving of control and data cells, as detailed in Sec. 7.

7 Uber

Fig. 9 gives a schematic for *Uber*. Typically, techniques rely on bypassing [17] or speculating router pipelines [20]. Set aside the additional hardware, these techniques complicate circuit timing closure [21], [23]. As a simpler alternative, *Uber* combines NoC ports. Although combining ports is a well-known practice [8], *Uber* is a novel application on NoCs. Unlike traditional interconnects, where latency is mainly determined by router chip IO [8], the advent of NoCs gives a good opportunity to rethink the role of buffers. Thus, seeing the buffers as the compiler and the router as the processor, *Uber* is a MIPS [12] equivalent for NoCs.

In an effort to tune the configuration developed in the previous sections into a realistic organization, there are two main steps to take. First, edges implement input in place of output queueing. Extensive evaluations suggest that results are hardly affected on average. Second, buffers are finite, and credit backpressure is added from the switch to the edges and from the edges to the cores. In turn, limited buffers imply deadlocks from protocol dependencies [8]. This paper follows the traditional approach to

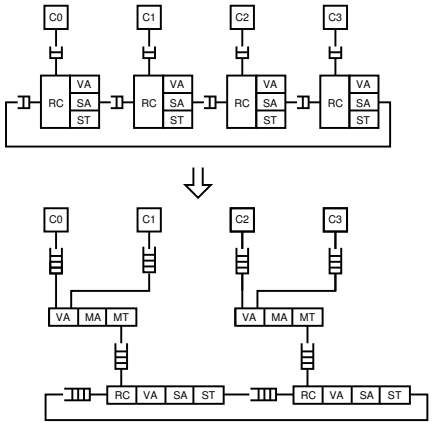

Fig. 9. *Uber* schematic

separate control messages (request, forward) and cache blocks (response) in virtual networks. Moreover, virtual networks are strictly prioritized. Preliminary results suggest that virtualization roughly halves end message delay. This differentiation does not change our conclusions. On the contrary, virtualization is a form of QoS for a critical part of the common case. Nevertheless, virtualization indeed relaxes the criticality of HoL blocking by nearly doubling the slack for queueing.

Thus, this paper predicts a sharply tripolar division of meshes will emerge with hundreds-cores. *Uber* will have a capacity around 4 Tb/s, will be highly loaded, and will use buffers as a counterweight for simpler routers. Note, however, that even using a small mesh, the NoC bandwidth is likely to be more critical for a particular subset of benchmarks (such as bodytrack, canneal, and facesim) than the whole set. Given that benchmark speedup is hard to sustain [14], the assumption on linear scaling of bandwidth is aggressive. Dark silicon [26] could be a second reason. The second pole will have a capacity beyond 16 Tb/s, will be lightly utilized, and will use more complicated router pipelines. And, the third pole will comprise even larger configurations and embarrassingly low utilizations, much like today's NoCs. Bufferless routing [19], as well as SMART [16] is particularly applicable for this third pole.

Ongoing work will contribute comparisons between *Uber* and bufferless routing [19], particularly with respect to energy. Ongoing work also address optimizations for *Uber*. Updating cache block placement to better balance load by reducing correlation (spikes) is a novel constraint that previous works have not taken into consideration [27].

8 Conclusion

Judging from state-of-the-art benchmarks —where corresponding slack for queueing is large in contrast to latency— utilizing buffers will play key role in scaling to hundreds-cores. Although *Uber* seems to keep queueing within the nominal slack, more work is needed on optimizations to improve scalability. Future work will also elaborate further by looking more carefully, inside the benchmark boxes. CUDA workloads are also important to analyze [2]. Either way, by no means has this paper found the ultimate NoC. This paper has contributed an important simplification, as well as a better understanding of stress.

REFERENCES

[1] N. Abeyratne, R. Das, Q. Li, K. Sewell, B. Giridhar, R. Dreslinski, D. Blaauw, and T. Mudge. Scaling towards kilo-core processors with asymmetric high-radix topologies. In *HPCA*, 2013.

[2] A. Bakhoda, G. Yuan, W. Fung, H. Wong, and T. Aamodt. Analyzing cuda workloads using a detailed gpu simulator. In *ISPASS*, 2009.

[3] J. Balfour and W. Dally. Design tradeoffs for tiled cmp on-chip networks. In *ICS*, 2006.

[4] N. Barrow-Williams, C. Fensch, and S. Moore. A communication characterization of splash-2 and parsec. In *IISWC*, 2009.

[5] C. Bienia. Benchmarking modern multiprocessors. *Phd Thesis, Princeton*, 2011.

[6] N. Binkert, B. Beckmann, G. Black, S. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. Hower, R. Derek, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M. Hill, , and D. Woo. The gem5 simulator. In *ACM SIGARCH C. Arch. News*, 2011.

[7] A. Carpenter, J. Hu, J. Xu, M. Huang, and H. Wu. A case for globally shared-medium on-chip interconnect. In *ISCA*, 2011.

[8] W. Dally and B. Towles. *Principles and Practices of Interconnection Networks*. 2003.

[9] R. Das, S. Eachempati, A. Mishra, V. Narayanan, and C. Das. Design and evaluation of a hierarchical on-chip interconnect for next-generation cmps. In *HPCA*, 2009.

[10] N. Enright-Jerger, L.-S. Peh, and M. Lipasti. Circuit switched coherence. In *NOCS*, 2008.

[11] P. Gratz and S. Keckler. Realistic workload characterization and analysis for networks-on-chip design. In *CMP-MSI*, 2010.

[12] J. Hennessy. Vlsi processor architecture. *IEEE Trans. Computers*, 1984.

[13] R. Hesse, J. Nicholls, and N. Enright-Jerger. Fine-grained bandwidth adaptivity in networks-on-chip using bidirectional channels. In *NOCS*, 2012.

[14] M. Hill and M. Marty. Amdahl's law in the multicore era. *Computer*, 2008.

[15] M. Karol, M. Hluchyj, and S. Morgan. Input versus output queueing on a space-division packet switch. In *IEEE Trans. Communications*, 1988.

[16] T. Krishna, C.-H. Chen, W. Kwon, and L.-S. Peh. Breaking the on-chip latency barrier using SMART. In *HPCA*, 2013.

[17] A. Kumar, L.-S. Peh, P. Kundu, and N. Jha. Express virtual channels: Towards the ideal interconnection fabric. In *ISCA*, 2007.

[18] M. Martin, M. Hill, and D. Sorin. Why on-chip cache coherence is here to stay. In *Communications ACM*, 2012.

[19] T. Moscibroda and O. Mutlu. A case for bufferless routing in on-chip networks. In *ISCA*, 2009.

[20] R. Mullins, A. West, and S. Moore. Low-latency virtual-channel routers for on-chip networks. In *ISCA*, 2004.

[21] R. Mullins, A. West, and S. Moore. The design and implementation of a low-latency on-chip network. In *ASP-DAC*, 2006.

[22] N. Mckeown. Scheduling algorithms for input-queued cell switches. *PhD Thesis, Univ. California at Berkeley*, 1995.

[23] S. Park, T. Krishna, C.-H. O. Chen, B. Daya, A. Chandrakasan, and L.-S. Peh. Approaching the theoretical limits of a mesh noc with a 16-node chip prototype in 45nm soi. In *DAC*, 2012.

[24] G. Passas, M. Katevenis, and D. Pnevmatikatos. The combined input-output queued (CIOQ) crossbar architecture for high-radix on-chip switches. In *IEEE Micro*, 2015.

[25] V. Paxson and S. Floyd. Wide area traffic: the failure of poisson modeling. *IEEE/ACM Trans. Networking*, 1995.

[26] M. Taylor. Is dark silicon useful?: Harnessing the four horsemen of the coming dark silicon apocalypse. In *DAC*, 2012.

[27] A. Udipi, N. Muralimanohar, and R. Balasubramonian. Towards scalable, energy-efficient, bus-based on-chip networks. In *HPCA*, 2010.

[28] S. Woo, M. Ohara, E. Torrie, J. Singh, and A. Gupta. The splash-2 programs: Characterization and methodological considerations. In *ISCA*, 1995.